

Inferring repeat protein energetics from evolutionary information

Espada Rocío¹², Parra R. Gonzalo³, Mora Thierry⁴, Walczak Aleksandra M.⁵, Ferreiro Diego U.^{12*}

¹ Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Química Biológica. Buenos Aires, Argentina.

² CONICET - Universidad de Buenos Aires. Instituto de Química Biológica de la Facultad de Ciencias Exactas y Naturales (IQUIBICEN). Buenos Aires, Argentina.

³ Quantitative and Computational Biology Group, Max Planck Institute for Biophysical Chemistry, Goettingen, Germany

⁴ Laboratoire de physique statistique, Ecole Normale Supérieure, CNRS and UPMC, 75005 Paris, France

⁵ CNRS and Laboratoire de Physique Théorique, Ecole Normale Supérieure, Paris, France

* Corresponding author Ferreiro Diego U., mail: ferreiro@qb.fcen.uba.ar.

February 2017

Abstract

Natural protein sequences contain a record of their history. A common constraint in a given protein family is the ability to fold to specific structures, and it has been shown possible to infer the main native ensemble by analyzing covariations in extant sequences. Still, many natural proteins that fold into the same structural topology show different stabilization energies, and these are often related to their physiological behavior. We propose a description for the energetic variation given by sequence modifications in repeat proteins, systems for which the overall problem is simplified by their inherent symmetry. We explicitly account for single amino acid and pair-wise interactions and treat higher order correlations with a single term. We show that the resulting evolutionary field can be interpreted with structural detail. We trace the variations in the energetic scores of natural proteins and relate them to their experimental characterization. The resulting energetic evolutionary field allows the prediction of the folding free energy change for several mutants, and can be used to generate synthetic sequences that are statistically indistinguishable from the natural counterparts.

Keywords: repeat proteins, coevolution analysis, ankyrin repeat, leucine-rich repeat, tetratricopeptide repeat, ΔG prediction.

Introduction

Repeat proteins are composed of tandem repetitions of similar structural motifs of about 20 to 40 amino acids. Under appropriate conditions, these polymers fold into elongated, non-globular structures (Fig. 1). It is apparent that the overall architecture is stabilized mainly by short range interactions, in contrast to most globular protein domains that usually adopt very intricate topologies [1]. In their natural context, repeat proteins are frequently found mediating protein-protein interactions, with a specificity rivaling that of antibodies [2, 3, 4]. Given their structural simplicity and potential technological applications, repeat-

proteins are a prime target for protein design, with very successful examples for a variety of topologies [5, 6, 7]. Most of the current design strategies target the creation of rigid native structures with desired folds that, although beautiful, often lose biological functionality [8]. It is becoming clear that the population of ‘excited states’ is crucial for protein function [9], and thus tackling energetic inhomogeneities in protein structures may be crucial for understanding how biological activities emerge [10]. The challenge thus relies in finding an appropriate description for the ‘energy’ of each system, a daunting task for large molecular objects such as natural proteins.

In principle, the natural variations observed for pro-

teins of the same family must contain information about the sequence-structure mapping. A simple model that just takes into account the frequency of each amino acid in each position is insufficient to capture collective effects, yet, for some architectures it is surprisingly good for the synthesis of non-natural repeat-proteins by ‘consensus’ design [11, 12, 13, 14]. It is apparent that in the case of repeat proteins the local signals play inordinately large roles in the energy distribution, just as expected from their topology [15] and hence, small heterogeneities can be propagated from the local repeat units to higher orders affecting the overall structure and dynamics [16, 17]. Thus, collective effects may be approximated as small perturbations to local potentials, simplifying the energetic description of complex natural systems [18].

In the last years new methods to analyze correlated mutations across a family of proteins have arisen (mfDCA[20], plmDCA [21, 22], Gremlin [23], EVFold [24] to name a few). The main hypothesis behind these methods is that biochemical changes produced by a point mutation should be compensated by other mutations (along evolutionary timescales) to maintain protein viability or function. These methods can also be used to disentangle relevant direct correlations from indirect ones. They are very successful at predicting spatial contacts and interactions for many protein topologies [25, 26, 27, 28, 29]. Nevertheless, these methods do not take into account the chemical nature of the amino acids, which can be codifying inhomogeneities in the energetic distribution that are crucial for the activity of repeat-proteins [30, 31]. On this basis, different approaches have been proposed recently to include chemical details in the correlation analyses [32], trying to predict folding stability [33], conformational heterogeneity [25, 34, 35] or the global effect on antibiotic resistance from sequences of β -lactamases [36, 37]. As many other tools, these were optimized to perform well on globular proteins, and their application to repeat proteins is not straightforward. Besides the point-mutation mechanism, repeat proteins are believed to evolve via du-

plication and rearrangement of repeats [38], resulting in an inherent symmetry which usually confounds sequence analyses [17]. Making use of this symmetry, we have previously proposed a specific version of mfDCA and plmDCA for repeat proteins [39]. In this work we develop an alternative ‘evolutionary field’ to approximate the biochemical properties of repeat proteins just from the analysis of natural sequences. We take advantage of the elongated and repetitive structure of these proteins (Fig. 1) to extract as much information as possible from the data, and apply the general ideas on three specific families, ankyrin repeats (ANK), leucine-rich repeats (LRR) and tetratricopeptide-like repeats (TPR).

Evolutionary energy for repeat proteins

To study the co-occurrence of mutations in a sequence alignment of a particular protein family, [40] proposed a Hamiltonian or energy expression which resembles a Potts model:

$$E(\vec{s}) = - \left[\sum_{i=1}^L h_i(a_i) + \sum_{i=1}^L \sum_{j=i}^L J_{ij}(a_i, b_j) \right] \quad (1)$$

where the set of $\{h_i(a_i)\}$ parameters, one for each amino acid in each position, accounts for a local propensity of having a specific residue on a particular site of the protein, and the set of $\{J_{ij}(a_i, b_j)\}$ indicates the strength of the ‘evolutionary’ interaction between each possible amino acid in every pair of positions along the protein. There are $q = 21$ possible values of a_i and b_j , one for each amino acid and one for the gaps included on the multiple sequence alignments. This expression is evaluated on a particular sequence on an alignment, and the summations go over the L columns of the alignment. A sequence is more favorable or more energetic if it gets lower values of $E(\vec{s})$. It can be expected that the population of sequences follows a Boltzmann distribution $P(\vec{s}) = \frac{1}{Z} e^{-E(\vec{s})}$ [41]. The parameters are thus fitted to reproduce the frequencies of occurrence

of each amino acid in each position ($f_i(a_i)$) and the joint frequencies of amino acids ($f_{ij}(a_i, b_j)$) in an alignment of natural sequences used as input:

$$f_i(a_i) = \sum_{a_k, k \neq i} P(\vec{s}) \quad (2)$$

$$f_{ij}(a_i, b_j) = \sum_{a_k, k \neq i, j} P(\vec{s}) \quad (3)$$

Nevertheless, for repeat proteins there is another feature we want to capture with an evolutionary energy: the high identity of amino acids constituting consecutive repeats, arisen by the repetitiveness of these families and probably a signature of their evolutionary mechanisms (Fig. 2).

Therefore, we propose the following model for repeat proteins:

$$E(\vec{s}) = - \left[\sum_{i=1}^L h_i(a_i) + \sum_{i=1}^L \sum_{j=i}^L J_{ij}(a_i, b_j) - \lambda_{Id}(\vec{s}) \right] \quad (4)$$

This expression is designed to be applied in sequences constituted by two repeats. λ_{Id} is a parameter that aims at reproducing the probabilities of the percentage of identity (%Id) between consecutive repeats in natural proteins (p_{id}). Basically, it accounts for higher order correlations not captured by the pairwise terms. For a given sequence we calculate the %Id of the adjacent repeats and sum the parameter λ_{Id} corresponding to that %Id value. When the correct parameters are obtained, this equation can be used to produce an ensemble of sequences consistent with the constraints ($f_i(a_i)$, $f_{ij}(a_i, b_j)$ and p_{id}). We work with pairs of repeats as it is the minimum unit that includes the

interaction between repeats and the possibility of measure sequence identity between consecutive repeats. In the following section we will show the convergence of the method and the relevant information that can be obtained from it. For further details about the procedure to assign values to the parameters, please refer to Methods section.

Results

Evolutionary Energy reproduces ensembles of sequences with natural frequencies and repeat protein characteristics.

We construct an alignment of pairs of repeats for each family: ANK (PFAM id PF00023, and final alignment of 20513 sequences of L=66 residues each), TPR (PFAM id PF00515, and final alignment of 10020 sequences of L=68 residues each) and LRR (PFAM id PF13516, and final alignment of 18839 sequences of L=48 residues each). See Methods for further details of construction. We measure $f_i(a_i)$, $f_{ij}(a_i, b_j)$ and p_{id} . Using a gradient descent procedure we obtain a set of parameters in equation 4 which are able to reproduce $f_i(a_i)$, $f_{ij}(a_i, b_j)$ and p_{id} . In principle, the number of parameters is large: Lq h_i parameters, $\frac{(Lq)^2 - Lq}{2} J_{ij}$ parameters and $\frac{L}{2} + 1 \lambda_{Id}$. For example, for pairs of ANK repeats this means 1386 h_i , 959805 J_{ij} and 34 λ_{Id} . To reduce the number of free parameters to fit we use a L_1 -regularization which fixes to zero those parameters which do not contribute significantly to fit the frequencies. This regularization allows us to set to exactly




Figure 1: Repeat proteins are elongated objects with internal symmetry. Representative structures of members of the repeat proteins families studied. On left, ankyrin repeat (PDB id:1N0R [19]), center tetratricopeptide-like repeats (PDB id:1NA0 [11]) and right leucine-rich repeats (PDB id:4IM6). The defined repeated unit is highlighted in orange.

zero between 85 and 91% of the J_{ij} parameters which, when they are free to vary, only reach small values (Suppl. Fig. 3). We bound the maximum error permitted in the frequency estimations to 0.02. Refer to Methods for more details.

In the three families studied, the parameters obtained allow us to generate ensembles of sequences which reproduce natural $f_i(a_i)$, $f_{ij}(a_i, b_j)$ and p_{id} (Fig. 2A). Notice that most frequencies are fitted with an error of an order of magnitude lower than the maximum bound imposed (Suppl. Fig. 2).




Figure 2: The proposed model fits the frequencies of amino acids and natural repeat-identities p_{id} . On A, we compare marginal frequencies $f_i(a_i)$ (red) and joint frequencies $f_{ij}(a_i, b_j)$ (black) on the natural ensemble of sequences (x-axis) and on the set of sequences generated by the model (y-axis). On B, we calculate the distribution of identity between repeats p_{id} for consecutive repeats (solid line), and for natural repeats which are not consecutive, i.e. they are not next to each other in the primary structure (dot lines). Consecutive repeats present a population with high identity between repeats that any pairs of repeats do not show. We compare the distribution produced by the model p_{id}^{model} (dots).

The p_{id} distributions are also very well reproduced (Fig. 2B). Not only the general shape, but also the populated long tail for highly similar repeats. It is not possible to obtain the same distribution only by fitting amino acid frequencies $f_i(a_i)$ and $f_{ij}(a_i, b_j)$, it is mandatory to explicitly include the p_{id} by including the parameters λ_{Id} (Suppl. Fig. 1), suggesting that higher order correlations must be accounted for describing these systems.

Evolutionary Energy distinguishes between proteins on a given family and other polypeptides

Once the set of parameters $\{h_i(a_i), J_{ij}(a_i, b_j), \lambda_{Id}\}$ is obtained, it can be used to score any sequence of L amino acids via equation 4. In this section we test if this measure is capable of distinguishing polypeptides that fold in a three dimensional structure similar to members of the repeat protein family from those that do not.

We calculate the distribution of energies of different sets of sequences (Fig. 3). The ensembles of natural sequences of each protein family used to learn the parameters have a unimodal distribution of energies centered around -100 (Fig. 3, dots). These distributions are clearly differentiated from the energies of random chains of residues (Fig. 3, red lines).

For a positive control we evaluate designed proteins which have been experimentally synthesized. For the ANK family, we consider the library of repeat sequences built by Plückthun’s laboratory [13] (blue dots line, Fig. 3A). This library was constructed by fixing on each repeat 26 positions out of 33 to the most frequent residue in the multiple sequence alignment. This resulted in a set of sequences that have small variations with respect to the ANK consensus (the sequence with the most frequent amino acid in each position). In our expression, they score a very low energy distribution, overlapping with the most negative tail of the distribution of natural sequences. It is notable that consensus designed ANK have been shown experimentally to be extremely stable. For the TPR family, consensus designed was done by Regan’s laboratory [12, 11]. All pairs of repeats synthesized have the same amino acid sequence, and its energy score is indicated by a green full square in Fig. 3B. Again, the designed sequence matches values at the most left side of the energy distribution of natural sequences, and coincidentally reports high folding stability. From it, other variants with few point mutations to improve binding to a specific ligand have been

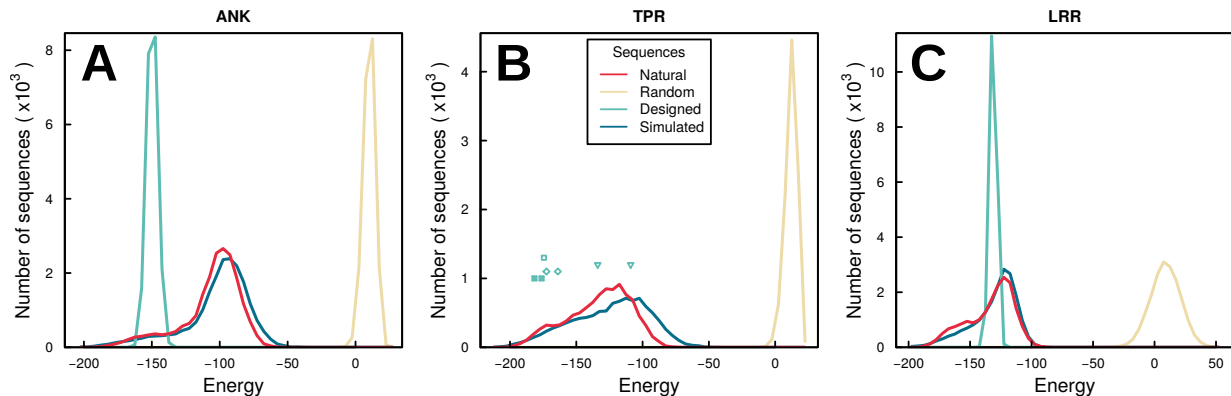


Figure 3: Energy score distribution for different ensembles of sequences. Red lines, natural sequences used to train the model on equation 4. Blue lines, sequences simulated by Monte Carlo under expression 4. In the three families, it overlaps with natural sequences, suggesting that simulated sequences imitate the natural ensemble. Yellow lines, strings of random amino acids used as negative control. They show that the energy distinguishes between polypeptides belonging to a protein family and other strings of amino acids. Green lines, squares, diamonds and triangles, energies for designed proteins.

synthesized. As shown in empty green squares [42] and diamonds [43] in Fig. 3B, they have higher energy, but still in the left most side of natural sequences distribution. Recently, a different design strategy was done [44]. Based on a non-repetitive protein, but similar to TPR fold, they put together various repetitions of the fold, using TPR loops to link them. They obtained a three-repeats protein whose pair of repeats energy are represented on triangles on Fig. 3B. This time, they match natural sequences distribution in higher values.

Finally, for the LRR family we contrast with the library of proteins designed by Plückthun’s group based on the consensus sequence [14]. The repetition they considered has 57 amino acids, which includes two types of repeats, one of 28 residues and the other one of 29. As the repeat we are using for LRR is 24 residues long, we aligned both definitions and evaluated the library removing the amino acids not matching our definition. Again, their scores form a narrow distribution, but this time it is not placed on the most favorable side of the natural sequences distribution (Fig. 3C). Coincidentally, selected species studied do not show such a high folding stability as the ANK library did.

With these parameters, we are able to generate an ensemble of sequences which are in agreement with the constraints used, via a Monte Carlo simulation (see Methods).

The distribution of energies of these simulated sequences matches the natural sequences energies distribution with remarkable accuracy. Moreover, we randomly choose 100 sequences from the natural ensemble and 100 sequences from the simulated one, perform a Smith-Waterman pairwise alignment all against all, calculate the pair similarity using BLOSUM62 matrix and used it as a distance method to plot a dendrogram of the sequences (Supl. Fig. 4). Both species appear interspersed, showing that it is not possible to distinguish a natural sequence from a constructed one. Also, we tested *familiarity* to the ANK family as defined in [45] and found overlapping distributions for both species (Supl. Fig. 5). Therefore, simulated sequences represent possible variants to natural repeats. The wide distribution of natural proteins suggests that it should be possible to engineer sequences with more variable repeats, more dissimilar among neighbors and to the consensus than the ones published up to date.

Low evolutionary energy sequences have similar repeats

Are there any invariant properties shared by low energy sequences? Given that repeat-proteins may evolve by other mechanisms besides point substitutions, we analyze if low energy sequences are constituted by highly similar repeats and if they are close to consensus sequences.

On Fig. 4A we show the relation between the %Id between the repeats and the energy of the sequence. It is evident that low energy sequences are constructed by pairs of highly similar repeats. This could be a transitive effect: if low energy sequences are very similar to the consensus sequence, and the consensus sequence is formed by two identical repeats, we would be seeing that more similarity between repeats causes lower energies. We can see that it is not the case (Fig. 4B). We plot the %Id to the consensus against the energy of each sequence. The consensus was calculated with the most frequent amino acid in each position on sequences used as input. We can see that there is no evident correlation between the energy and the similarity to the consensus. Thus, low energy sequences that differ from the consensus one may be constructed. Also, there are no sequences which get a high %Id to the consensus. We conclude that there are different repeats which have low energies within a protein family, and not only the consensus sequence.




Figure 4: Most favorable simulated sequences have very similar repeats, yet they are different to the consensus repeat. On A, we plot the energy vs. the identity between the repeats that constitute the sequence. Even though the deviation is large, most stable sequences tend to have more similar repeats. On B, we plot the energies of simulated sequences vs the identity to the consensus of the family. In all cases, the identity to the consensus is small and uncorrelated to the energy, indicating that sequences which differ significantly from the consensus can be stable variants of the family.

Evolutionary Energy and folding stability change upon point mutations

Consensus designed ANK proteins are very stable upon chemical and thermal denaturation [13], and, as shown in Fig. 3 also score a very low evolutionary energy according to equation 4. Can we quantify the relationship between the stability and the evolutionary energy?

A potential test can be performed by comparing to experiments in which the effect of point mutations was evaluated. These incorporate one, two or three point mutations in natural proteins, and characterize the unfolding free energy ΔG of the wildtype and the mutated variant. A higher ΔG reports a more stable protein. We compare the change in the ΔG between the mutated and the wild-type protein ($\Delta\Delta G$), and the difference of energy for their sequences according to equation 4.

Although the energy expression is learned for pairs of repeats, we can easily extend it to an array of repeats making use of the elongated structure of repeat proteins in which only adjacent repeats interact. From our expression we have parameters assigned to intra-repeat positions (h_i with $i = 1 \dots \frac{L}{2}$ and J_{ij} with $i, j = 1 \dots \frac{L}{2}$), and inter-repeat interactions (J_{ij} with $i = 1 \dots \frac{L}{2}$ and $j = \frac{L}{2} + 1 \dots L$, and λ_{Id}). Then for each repeat we can assign an internal energy $\sum_{i=1}^{L/2} h_i(a_i) + \sum_{i=1}^{L/2} \sum_{j>i}^{L/2} J_{ij}(a_i, b_j)$ and a interaction energy $\sum_{i=1}^{L/2} \sum_{j=L/2+1}^L J_{ij}(a_i, b_j) + \lambda_{Id}$, which of course depends on the amino acids constituting each repeat.

On Fig. 5A, we show the comparison between $\Delta\Delta G$ and the evolutionary energy calculated using Eq. 4, done for three different ANK proteins: I κ B α [46, 47], Notch [48] and p16 [49]. It should be noted that different experimental techniques return different values for ΔG for the same protein, non overlapping within experimental error, pointing that other factors contribute to the experimental quantification of $\Delta\Delta G$. A linear fit returns $R^2 \approx 0.49$. Nevertheless, from 152 mutations we analyzed, 114 (75 %) are predicted favorable when the mutation stabilized the folding of the structure, and unfavorable when they

have also been measured to destabilize. The predictions that deviated the most are mutations in Notch from Serine to Proline, which is a structural disruptor, and were not considered in the linear fit.

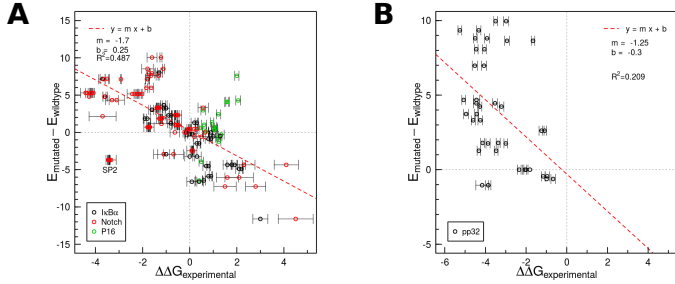


Figure 5: Variation of energy score as a predictor of the folding stability upon point mutations. We compare difference in unfolding ΔG between a wildtype protein and a mutated variant (x-axis) and the change in energy according to Eq. 4. Error bars indicate the experimental standard deviation. On A, for proteins belonging to ANK family, and on B for LRR.

On Fig. 5B, we show reported mutations on pp32 [50], a protein belonging to LRR family. Again, measurements with different methods report different values of $\Delta\Delta G$. The linear fit returns a poor $R^2 \approx 0.21$, but 30 (75%) mutations are both predicted and reported unstabilizing.

A similar comparison was performed by [33] for small globular proteins with an expression related to Eq. 1. To reduce the number of interaction parameters $J_{ij}(a_i, b_j)$ they explicitly used structural information and set to zero all interactions between positions which are not in contact in the native structure. In contrast, we use a L_1 -regularization to fix to zero those parameters which do not contribute significantly to the fitting process and obtain $J_{ij}(a_i, b_j) = 0$ and $J_{ij}(a_i, b_j) \neq 0$ in all pairs of positions, regardless they are supposed to be in contact or not in the 3D structure.

Interaction parameters are related to the structure and the sequence symmetry

Are the obtained parameters related to structural properties of these proteins? Local fields, $h_i(a_i)$, should account for the local propensity of each amino acid in each position, and therefore are expected to be related to $f_i(a_i)$. Fig. 6A shows that the inferred $h_i(a_i)$ parameters are dif-

ferent from the initial condition $-\ln(f_i(a_i))$ for the ANK family; that is, the values obtained for the parameters that account for higher order correlations are relevant. In red we highlight the points related to the consensus amino acid in each position. All of these residues have a strong local field associated to them, justifying why the construction of sequences with these amino acids results in foldable proteins. We also show a contact map of two ANK repeats (PDB id: 1N0R) on Fig. 6B: gray background indicates that the two positions given by x and y axis are in contact in the native structure, and white that they are not. On the upper triangle of the figure and in blue crosses, we mark the positions involved in the highest J_{ij} parameters, i.e. those which imply higher coupling. A darker blue indicates that there are more J_{ij} (more combinations of amino acids) between those positions. Most of the highest J_{ij} match a pair of positions in contact in the 3D structure, or two which correspond to the same residue in the adjacent repeat patterns, i.e. i -th position in the first repeat and position $j=i+33$ in the second repeat. In red crosses we show the lowest J_{ij} , that mark a negative constraint. Again, a darker red means that there are more J_{ij} with low values between those positions. It is apparent that these also involve mostly residues in contact, but shows that other regions are responsible for negative design.

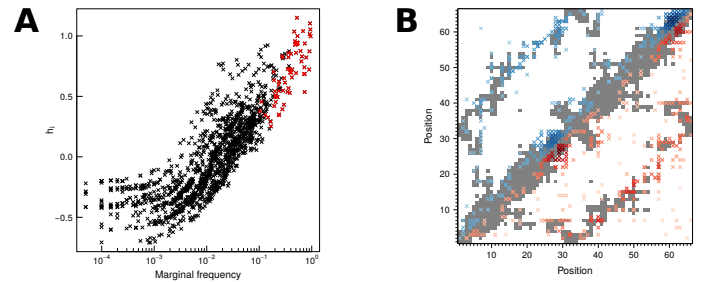


Figure 6: For the ANK family, on panel A we compare the parameters $h_i(a_i)$ to the marginal frequencies. The site-independent model (and initial condition) states that $h_i(a_i) = \ln(f_i(a_i))$. For the final model, this relation is tuned by the higher order correlations. On red, the parameters associated with the most common amino acid in each position are highlighted. On panel B, we compare the contact map of a pair of repeats of 1N0R (gray shadow) and the highest (blue) and lowest (red) $J_{ij}(a_i, b_j)$ parameters. The color scale indicates how many parameters involves the two positions (due to different sets of amino acids). Most extreme values fall into residues in contact or in the equivalent position of a repeat.

Discussion

We propose a statistical model to account for fine details of the energy distribution in families of repeat proteins using only the sequences of amino acids. The model consists of a generalization of a Potts model to account for the local and pair-wise interactions and an extra term that includes higher order correlations, accounting for the similarity between consecutive repeats. The model is constrained by evolutionary characteristics of the families of proteins: we measure the frequencies of amino acids, co-occurrence of amino acids and the identity between repeats in extant natural proteins. To statistically define these quantities it is necessary to have a large set of sequences, which we showed are currently available for several repeat-protein families [39]. No information about the native folded conformation is required. The computation of the evolutionary energy field is computationally demanding, mostly due to long times spent in rigorous Monte Carlo simulations, but once the fitting is done the parameters can be used to score individual sequences fast and easily.

We studied three popular repeat protein families: ANK, TPR and LRR. After pre-processing of the alignments, we had enough sequences (≈ 20500 , 10000 and 18800 respectively) to fit the model to pairs of repeats of each family. We scored the *evolutionary energy* of all natural sequences in PFAM, and it allowed us to clearly distinguish between natural proteins and random sequences of amino acids: the first have energy values < -50 and show a large spread while all random sequences have energy values ≈ 0 . We evaluated designed repeat proteins which have been shown to fold and found that they score within the natural sequences distribution of energies. For the ANK and TPR family, these designed proteins have been shown to be highly stable upon thermal and chemical denaturation and, coincidentally, they are located at the most favorable side of the energy distribution of natural proteins, suggesting that the evolutionary energy score can be related to folding stability.

The energetic model can be used in Monte Carlo sim-

ulations to generate sequences that agree with the natural constraints of a given protein family. This ensemble of simulated sequences matches the amino acid frequencies, the identity between repeats and also the energy distribution of natural proteins. We found this set of simulated sequences is statistically indistinguishable from natural counterparts. Thus, the proposed model can be used as a tool to design repeat-protein sequences that have all the natural characteristics evaluated to date. Moreover, the stability change upon single point mutation can be well predicted by the model. For both the simulated sequences and for natural counterparts, we found that the similarity between consecutive repeats correlates with lower energy values, and that these are not necessarily similar to the consensus sequence of the family, pointing out that duplication of stretches of sequences may well be an important factor in the evolution of these systems [51].

The existence of a simple and reliable energy function to score the ‘evolutionary energy’ of repeat-proteins can be used to trace the biological forces that acted upon their history, and to explore to which extent these conflict with the physical necessities of the systems [52]. Mapping the energy inhomogeneities along the repeat-arrays may allow us to infer the population of excited states in these proteins, many of which have been related to their physiological mechanisms.

Methods

Sequence alignments

Sequences of repeats were obtained from PFAM 27.0 [53]. These sequences usually have misdetected initial and final residues. We completed these positions with the amino acids present on the actual proteins. This leads to a reduction on the number of gaps in our alignments, which usually derives into noisy predictions in correlation analyses [33]. After, we created the alignment of pairs of repeats, joining sequences of repeats which are consecutive in a natural protein. Finally, we removed insertions from

the alignments by deleting positions which have gaps in more than 80% of the sequences in the alignment.

Frequency calculations

Our model fits the occurrence of amino acids in every position, which we call the marginal frequency of residue a_i at position i of the alignment and denote $f_i(a_i)$, and the joint occurrence of two amino acids a_i and b_j simultaneously at two different positions of the alignment, $f_{ij}(a_i, b_j)$. To avoid biases by the overrepresentation of some proteins in the database, we used CD-HIT [54] to choose representative sequences which differ between them in more than 90% of identity percentage. Finally, we computed by counting the $f_i(a_i)$ and $f_{ij}(a_i, b_j)$, and divided by the total number of sequences.

p_{id} calculations

From the same alignment explained in *Frequencies calculations*, for a sequence which has L residues constituting two consecutive repeats, the %Id between the repeats is the number of amino acids in positions i and $i + \frac{L}{2}$, for $i = 1 \dots \frac{L}{2}$ which are exactly the same. Gaps are treated as an amino acid. Once we have the values for all sequences in an alignment, we define p_{id} as the proportion of sequences within the alignment with the same %Id between repeats.

Construction of an ensemble of sequences in agreement with a energy equation

Given a set of parameters $h_i, J_{ij}, \lambda_{Id}$ and Eq. 4, we use a Monte Carlo procedure and the Metropolis criterion to generate an ensemble of N sequences of length L each. We initiate with a random string of L residues. At each step, we produce a point mutation in any position. If this mutation is favorable, i.e. the energy is lower than that of the original sequence, we accept the mutation. If not, we accept the mutation with a probability of $e^{-\Delta E}$, where ΔE is the difference of energy between the original and

the mutated sequence. When accepted, the mutated sequence is used as the original one for next step. We add one sequence to our final ensemble every t steps (we used $t=1000$).

Learning the parameters for the model

Our model is proposed to reproduce $f_i(a_i)$, $f_{ij}(a_i, b_j)$ and p_{id} from the alignment of natural sequences. To learn the set of parameters $h_i, J_{ij}, \lambda_{Id}$ which reproduce them, we used a gradient descent procedure. In each step, an ensemble of $N=80000$ sample sequences was produced via Monte Carlo using as energy the expression 4 and the trial parameters. We measured its marginal, joint frequencies and p_{id} and updated the local parameters according to:

$$h_i^{t+1} \leftarrow h_i^t - \epsilon_s [f_i(a_i) - f_i^{model}(a_i)] \quad (5)$$

As the number of parameters for coupling is large ($= 21^2 L^2$), we used a regularization L_1 to force to 0 those parameters which are not contributing significantly to the modeled frequencies. Then, we update these parameters by:

$$J_{ij}^{t+1} \leftarrow \begin{cases} 0 & \text{if } J_{ij}^t = 0 \text{ and } \Delta < \gamma \\ \epsilon_j [-\gamma \text{sign}(\Delta)] & \text{if } J_{ij}^t = 0 \text{ and } \Delta > \gamma \\ J_{ij}^t + \epsilon_j [\Delta - \gamma \text{sign}(J_{ij}^t)] & \text{if } [J_{ij}^t + \epsilon_j (\Delta - \gamma \text{sign}(J_{ij}^t))] \cdot J_{ij}^t > 0 \\ 0 & \text{if } [J_{ij}^t + \epsilon_j (\Delta - \gamma \text{sign}(J_{ij}^t))] \cdot J_{ij}^t < 0 \end{cases} \quad (6)$$

where $\Delta = f_i(a_i) - f_i^{model}(a_i)$. Finally, the parameters λ_{Id} are updated according to:

$$\lambda_{Id}^{t+1} \leftarrow \lambda_{Id}^t + \epsilon_{ID} [p_{id}(\%Id) - p_{id}^{model}(\%Id)] \quad (7)$$

We iterated until the maximum difference between the predicted frequencies and the natural sequences was below 0.02. The code was written in C++ and is available at GitHub.

Acknowledgments

This work was supported by Universidad de Buenos Aires and Consejo Nacional de Investigaciones Científicas y Técnicas CONICET (ANPCyT grant PICT2012-1647 to R.E.

and D.U.F.). Work in Paris was supported by grant ERC-StG n. 306312. R.G.P. is a long-term EMBO postdoctoral fellow. R.E. and R.G.P. would like to thank C.F.K.

References

- [1] Ewan RG Main, Alan R Lowe, Simon GJ Mochrie, Sophie E Jackson, and Lynne Regan. A recurring theme in protein engineering: the design, stability and folding of repeat proteins. *Current opinion in structural biology*, 15(4):464–471, 2005.
- [2] Steven G Sedgwick and Stephen J Smerdon. The ankyrin repeat: a diversity of interactions on a common structural framework. *Trends in biochemical sciences*, 24(8):311–316, 1999.
- [3] Bostjan Kobe and Andrey V Kajava. The leucine-rich repeat as a protein recognition motif. *Current opinion in structural biology*, 11(6):725–732, 2001.
- [4] Gregory L Blatch and Michael Lässle. The tetratricopeptide repeat: a structural motif mediating protein-protein interactions. *Bioessays*, 21(11):932–939, 1999.
- [5] Pamela JE Rowling, Elin M Sivertsson, Albert Perez-Riba, Ewan RG Main, and Laura S Itzhaki. Dissecting and reprogramming the folding and assembly of tandem-repeat proteins. *Biochemical Society Transactions*, 43(5):881–888, 2015.
- [6] TJ Brunette, Fabio Parmeggiani, Po-Ssu Huang, Gira Bhabha, Damian C Ekiert, Susan E Tsutakawa, Greg L Hura, John A Tainer, and David Baker. Exploring the repeat protein universe through computational protein design. *Nature*, 2015.
- [7] Agathe Urvoas, Asma Guellouz, Marie Valerio-Lepiniec, Marc Graille, Dominique Durand, Danielle C Desravines, Herman van Tilbeurgh, Michel Desmadril, and Philippe Minard. Design, production and molecular structure of a new family of artificial alpha-helical repeat proteins (α rep) based on thermostable heat-like repeats. *Journal of molecular biology*, 404(2):307–327, 2010.
- [8] Rocío Espada, Ignacio E Sánchez, and Diego U Ferreira. Detailing protein landscapes under pressure. *Biophysical Journal*, 111(11):2339–2341, 2016.
- [9] Hans Frauenfelder. Function and dynamics of myoglobin. *Annals of the New York Academy of Sciences*, 504:151–167, 1987.
- [10] Diego U Ferreira, Elizabeth A Komives, and Peter G Wolynes. Frustration in biomolecules. *Quarterly reviews of biophysics*, 47(04):285–363, 2014.
- [11] Ewan RG Main, Yong Xiong, Melanie J Cocco, Luca D’Andrea, and Lynne Regan. Design of stable α -helical arrays from an idealized tpr motif. *Structure*, 11(5):497–508, 2003.
- [12] Tommi Kajander, Aitziber L Cortajarena, Simon Mochrie, and Lynne Regan. Structure and stability of designed tpr protein superhelices: unusual crystal packing and implications for natural tpr proteins. *Acta Crystallographica Section D: Biological Crystallography*, 63(7):800–811, 2007.
- [13] H Kaspar Binz, Michael T Stumpp, Patrik Forrer, Patrick Amstutz, and Andreas Plückthun. Designing repeat proteins: well-expressed, soluble and stable proteins from combinatorial libraries of consensus ankyrin repeat proteins. *Journal of molecular biology*, 332(2):489–503, 2003.
- [14] Michael T Stumpp, Patrik Forrer, H Kaspar Binz, and Andreas Plückthun. Designing repeat proteins: modular leucine-rich repeat protein libraries based on the mammalian ribonuclease inhibitor family. *Journal of molecular biology*, 332(2):471–487, 2003.
- [15] Diego U Ferreira, Aleksandra M Walczak, Elizabeth A Komives, and Peter G Wolynes. The energy landscapes of repeat-containing proteins: topol-

- p>ogy, cooperativity, and the folding funnels of one-dimensional architectures.
- PLoS computational biology*
- , 4(5):e1000070, 2008.
- [16] R Gonzalo Parra, Rocío Espada, Ignacio E Sánchez, Manfred J Sippl, and Diego U Ferreiro. Detecting repetitions and periodicities in proteins by tiling the structural space. *The Journal of Physical Chemistry B*, 117(42):12887–12897, 2013.
- [17] Rocío Espada, R Gonzalo Parra, Manfred J Sippl, Thierry Mora, Aleksandra M Walczak, and Diego U Ferreiro. Repeat proteins challenge the concept of structural domains. *Biochemical Society Transactions*, 43(5):844–849, 2015.
- [18] Hans Frauenfelder, Benjamin H McMahon, and Paul W Fenimore. Myoglobin: the hydrogen atom of biology and a paradigm of complexity. *Proceedings of the National Academy of Sciences*, 100(15):8615–8617, 2003.
- [19] Leila K Mosavi, Daniel L Minor, and Zheng-yu Peng. Consensus-derived structural determinants of the ankyrin repeat motif. *Proceedings of the National Academy of Sciences*, 99(25):16029–16034, 2002.
- [20] Faruck Morcos, Andrea Pagnani, Bryan Lunt, Arianna Bertolino, Debora S Marks, Chris Sander, Riccardo Zecchina, José N Onuchic, Terence Hwa, and Martin Weigt. Direct-coupling analysis of residue co-evolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*, 108(49):E1293–E1301, 2011.
- [21] Magnus Ekeberg, Cecilia Lövkvist, Yueheng Lan, Martin Weigt, and Erik Aurell. Improved contact prediction in proteins: using pseudolikelihoods to infer potts models. *Physical Review E*, 87(1):012707, 2013.
- [22] Magnus Ekeberg, Tuomo Hartonen, and Erik Aurell. Fast pseudolikelihood maximization for direct-coupling analysis of protein structure from many homologous amino-acid sequences. *Journal of Computational Physics*, 276:341–356, 2014.
- [23] Sivaraman Balakrishnan, Hetunandan Kamisetty, Jaime G Carbonell, Su-In Lee, and Christopher James Langmead. Learning generative models for protein fold families. *Proteins: Structure, Function, and Bioinformatics*, 79(4):1061–1078, 2011.
- [24] Thomas A Hopf, Charlotta PI Schärfe, João PGLM Rodrigues, Anna G Green, Oliver Kohlbacher, Chris Sander, Alexandre MJJ Bonvin, and Debora S Marks. Sequence co-evolution gives 3d contacts and structures of protein complexes. *Elife*, 3:e03430, 2014.
- [25] Faruck Morcos, Biman Jana, Terence Hwa, and José N Onuchic. Coevolutionary signals across protein lineages help capture multiple protein conformations. *Proceedings of the National Academy of Sciences*, 110(51):20533–20538, 2013.
- [26] Ryan R Cheng, Faruck Morcos, Herbert Levine, and José N Onuchic. Toward rationally redesigning bacterial two-component signaling systems using coevolutionary information. *Proceedings of the National Academy of Sciences*, 111(5):E563–E571, 2014.
- [27] Joanna I Sułkowska, Faruck Morcos, Martin Weigt, Terence Hwa, and José N Onuchic. Genomics-aided structure prediction. *Proceedings of the National Academy of Sciences*, 109(26):10340–10345, 2012.
- [28] Sergey Ovchinnikov, Hetunandan Kamisetty, and David Baker. Robust and accurate prediction of residue–residue interactions across protein interfaces using evolutionary information. *Elife*, 3:e02030, 2014.
- [29] Ryan R Cheng, Mohit Raghunathan, Jeffrey K Noel, and José N Onuchic. Constructing sequence-dependent protein models using coevolutionary information. *Protein Science*, 25(1):111–122, 2016.

- [30] Christian Zahnd, Emanuel Wyler, Jochen M Schwenk, Daniel Steiner, Michael C Lawrence, Neil M McKern, Frédéric Pecorari, Colin W Ward, Thomas O Joos, and Andreas Plückthun. A designed ankyrin repeat protein evolved to picomolar affinity to her2. *Journal of molecular biology*, 369(4):1015–1028, 2007.
- [31] Matthew J Cliff, Mark A Williams, John Brooke-Smith, David Barford, and John E Ladbury. Molecular recognition via coupled folding and binding in a tpr domain. *Journal of molecular biology*, 346(3):717–732, 2005.
- [32] Ronald M Levy, Allan Haldane, and William F Flynn. Potts hamiltonian models of protein co-variation, free energy landscapes, and evolutionary fitness. *Current Opinion in Structural Biology*, 43:55–62, 2017.
- [33] A Contini and G Tiana. A many-body term improves the accuracy of effective potentials based on protein coevolutionary data. *The Journal of chemical physics*, 143(2):025103, 2015.
- [34] Ludovico Sutto, Simone Marsili, Alfonso Valencia, and Francesco Luigi Gervasio. From residue co-evolution to protein conformational ensembles and functional dynamics. *Proceedings of the National Academy of Sciences*, 112(44):13567–13572, 2015.
- [35] Allan Haldane, William F Flynn, Peng He, RSK Vijayan, and Ronald M Levy. Structural propensities of kinase family proteins from a potts model of residue co-variation. *Protein Science*, 2016.
- [36] Matteo Figliuzzi, Hervé Jacquier, Alexander Schug, Oliver Tenaillon, and Martin Weigt. Coevolutionary landscape inference and the context-dependence of mutations in beta-lactamase tem-1. *Molecular biology and evolution*, 2015. doi: 10.1093/molbev/msv211.
- [37] Thomas A Hopf, John B Ingraham, Frank J Poelwijk, Michael Springer, Chris Sander, and Debora S Marks. Quantification of the effect of mutations using a global probability model of natural sequence variation. *arXiv preprint arXiv:1510.04612*, 2015.
- [38] Andreas Schüller and Erich Bornberg-Bauer. Evolution of protein domain repeats in metazoa. *Molecular Biology and Evolution*, 2016.
- [39] Rocío Espada, R Gonzalo Parra, Thierry Mora, Aleksandra M Walczak, and Diego U Ferreira. Capturing coevolutionary signals in repeat proteins. *BMC bioinformatics*, 16(1):207, 2015.
- [40] Martin Weigt, Robert A White, Hendrik Szurmant, James A Hoch, and Terence Hwa. Identification of direct residue contacts in protein–protein interaction by message passing. *Proceedings of the National Academy of Sciences*, 106(1):67–72, 2009.
- [41] Alexei V Finkelstein, Azat Ya Badretdinov, and Alexander M Gutin. Why do protein architectures have boltzmann-like statistics? *Proteins: Structure, Function, and Bioinformatics*, 23(2):142–150, 1995.
- [42] Anne Marie Krachler, Amit Sharma, and Colin Kleanthous. Self-association of tpr domains: Lessons learned from a designed, consensus-based tpr oligomer. *Proteins: Structure, Function, and Bioinformatics*, 78(9):2131–2143, 2010.
- [43] Aitziber L Cortajarena, Jimin Wang, and Lynne Regan. Crystal structure of a designed tetratricopeptide repeat module in complex with its peptide ligand. *FEBS journal*, 277(4):1058–1066, 2010.
- [44] Hongbo Zhu, Edgardo Sepulveda, Marcus D Hartmann, Manjunatha Kogenaru, Astrid Ursinus, Eva Sulz, Reinhard Albrecht, Murray Coles, Jörg Martin, and Andrei N Lupas. Origin of a folded repeat protein from an intrinsically disordered ancestor. *eLife*, 5:e16761, 2016.
- [45] Pablo Turjanski, R Gonzalo Parra, Rocío Espada, Verónica Becher, and Diego U Ferreira. Protein repeats from first principles. *Scientific reports*, 6, 2016.

- [46] Ingrid DeVries, Diego U Ferreiro, Ignacio E Sánchez, and Elizabeth A Komives. Folding kinetics of the cooperatively folded subdomain of the $\kappa\text{b}\alpha$ ankyrin repeat domain. *Journal of molecular biology*, 408(1): 163–176, 2011.
- [47] Diego U Ferreiro, Carla F Cervantes, Stephanie ME Truhlar, Samuel S Cho, Peter G Wolynes, and Elizabeth A Komives. Stabilizing $\kappa\text{b}\alpha$ by consensus design. *Journal of molecular biology*, 365(4):1201–1216, 2007.
- [48] Timothy O Street, Christina Marchetti Bradley, and Doug Barrick. An improved experimental system for determining small folding entropy changes resulting from proline to alanine substitutions. *Protein science*, 14(9):2429–2435, 2005.
- [49] Kit S Tang, Alan R Fersht, and Laura S Itzhaki. Sequential unfolding of ankyrin repeats in tumor suppressor p16. *Structure*, 11(1):67–73, 2003.
- [50] Thuy Phuong Dao, Ananya Majumdar, and Doug Barrick. Highly polarized c-terminal transition state of the leucine-rich repeat domain of pp32 is governed by local stability. *Proceedings of the National Academy of Sciences*, 112(18):E2298–E2306, 2015.
- [51] Åsa K Björklund, Diana Ekman, and Arne Elofsson. Expansion of protein domain repeats. *PLoS Comput Biol*, 2(8):e114, 2006.
- [52] Faruck Morcos, Nicholas P Schafer, Ryan R Cheng, José N Onuchic, and Peter G Wolynes. Coevolutionary information, protein folding landscapes, and the thermodynamics of natural selection. *Proceedings of the National Academy of Sciences*, 111(34):12408–12413, 2014.
- [53] Alex Bateman, Lachlan Coin, Richard Durbin, Robert D Finn, Volker Hollich, Sam Griffiths-Jones, Ajay Khanna, Mhairi Marshall, Simon Moxon, Erik LL Sonnhammer, et al. The pfam protein families database. *Nucleic acids research*, 32(suppl 1): D138–D141, 2004.
- [54] Weizhong Li, Lukasz Jaroszewski, and Adam Godzik. Tolerating some redundancy significantly speeds up clustering of large protein databases. *Bioinformatics*, 18(1):77–82, 2002.

Supplementary information




Figure 1: Distribution of % Id of sequences generated by model on equation 1 (main text) in red solid lines and by model 4 (main text) in black solid lines. In black dots, the natural sequences' distribution of %Id.




Figure 2: Histogram of frequency errors.




Figure 3: Change of J_{ij} parameters under regularization.

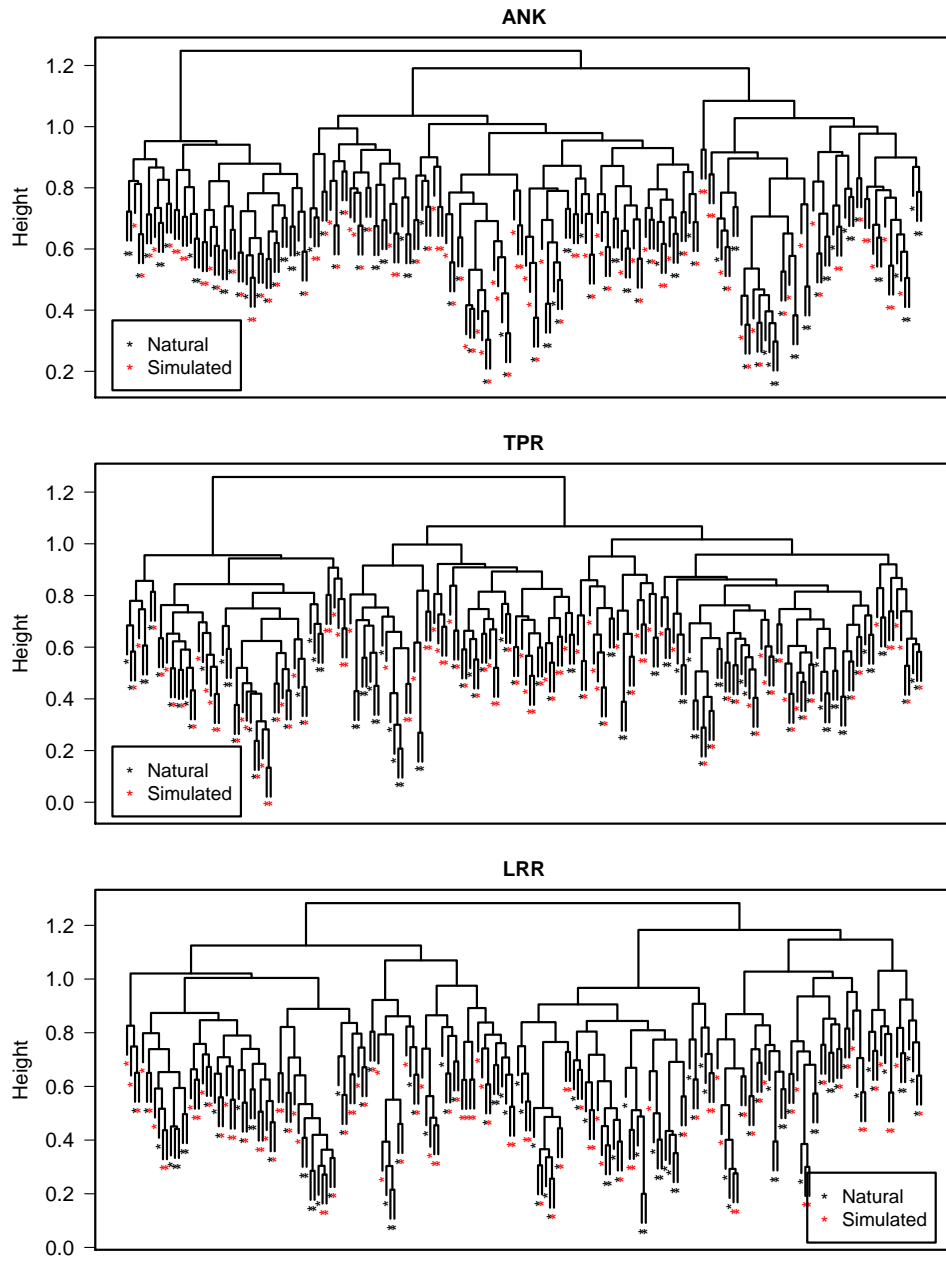


Figure 4: Dendrogram based on pair-wise similarity. Natural and simulated sequences are indistinguishable from pairwise similarity.

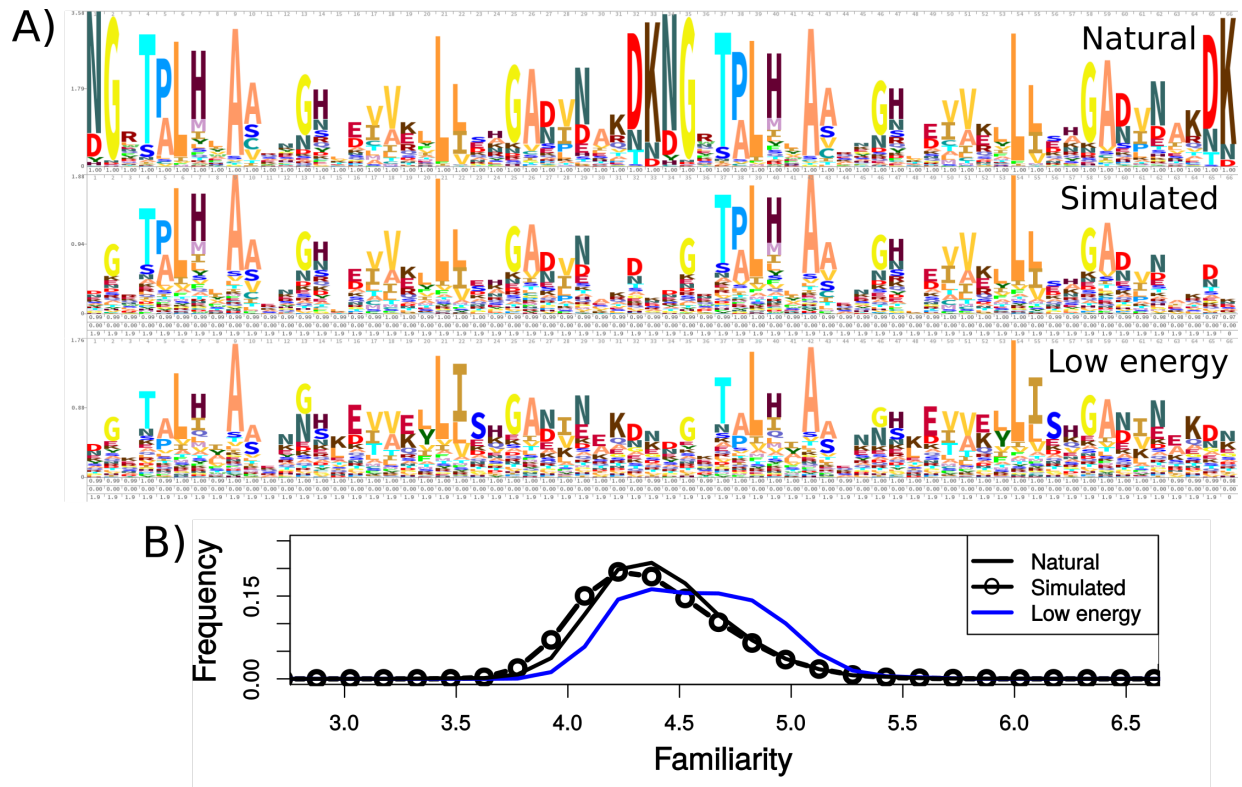


Figure 5: A) Logos for the MSA of natural pairs of ANK repeats (top), of simulated pairs of ANK repeats (center) and low energy pairs of simulated ANK repeats (bottom). B) Distribution of familiarity, as defined in Turjanski et al (2016) for the same sets of sequences. Simulated sequences reproduce the distribution of Natural proteins and are indistinguishable.

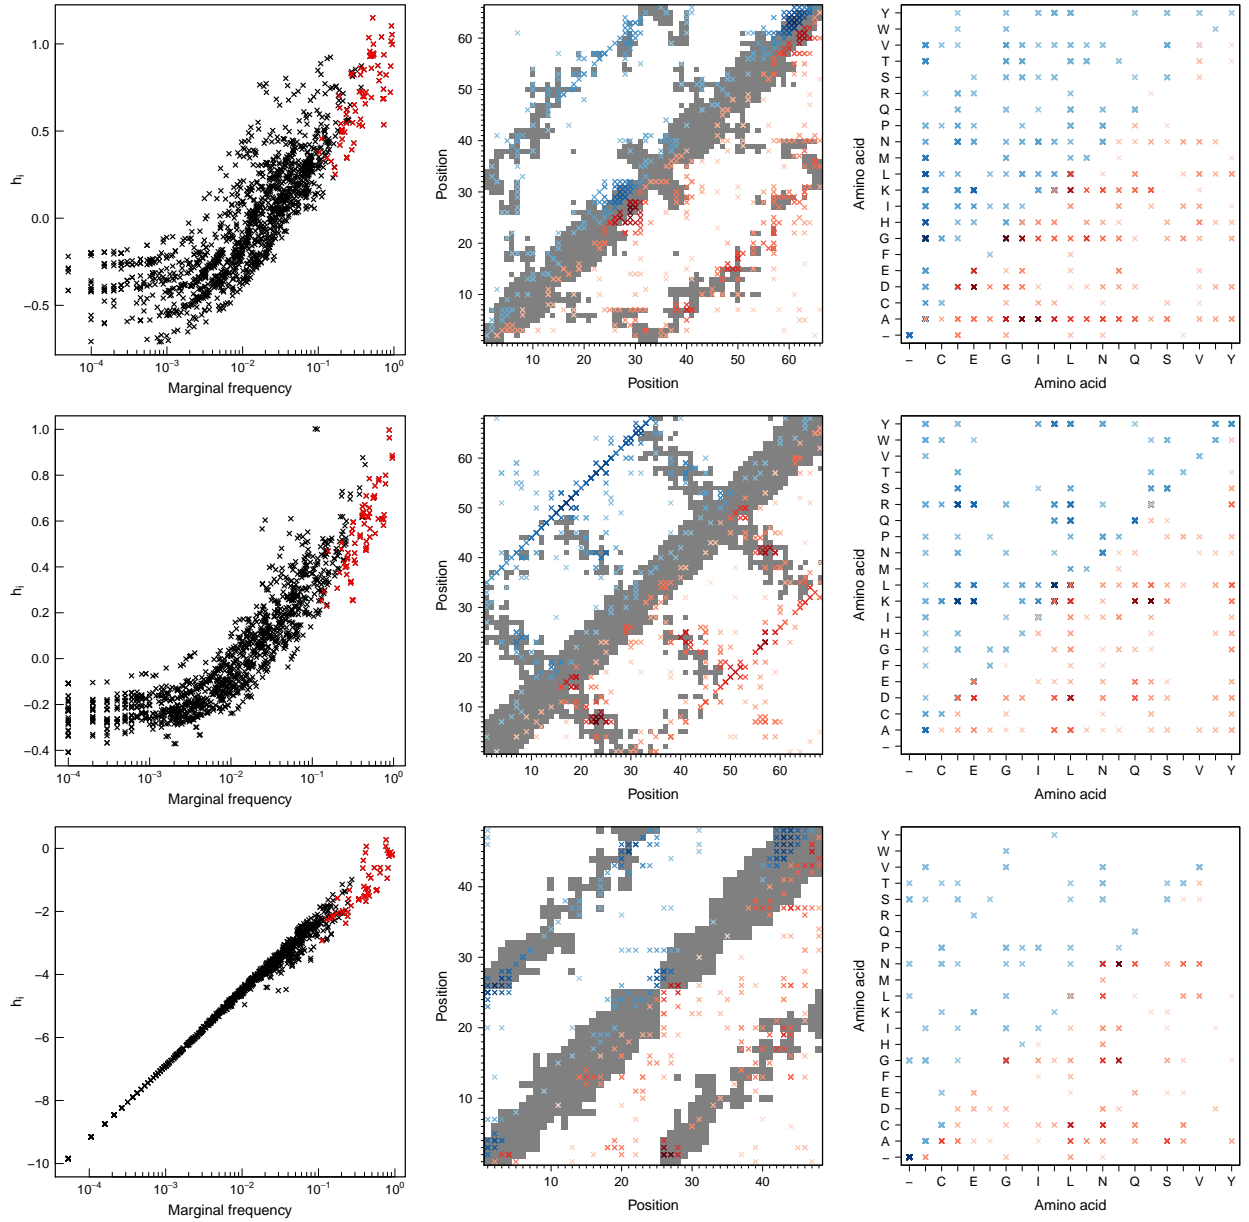


Figure 6: At left, comparison between the local field parameters $h_i(a_i)$ and the marginal frequencies $f_i(a_i)$. At center, contact map (grey indicates position in contact, white not in contact on the native structure). On blue, pairs of positions involved in highest $J_{ij}(a_i, b_j)$, red lowest $J_{ij}(a_i, b_j)$. At right, pairs of amino acids involved in the highest $J_{ij}(a_i, b_j)$ parameters (on blue) and in the lowest $J_{ij}(a_i, b_j)$ (on red). ANK at the top. TPR at the center. LRR at the bottom.